

# A mixed model approach for joint genetic analysis of alternatively spliced transcript isoforms using RNA-Seq data

Barbara Rakitsch<sup>1</sup>, Christoph Lippert<sup>2</sup>, Hande Topa<sup>3</sup>, Karsten Borgwardt<sup>1,4</sup>, Antti Honkela<sup>5</sup>, and Oliver Stegle<sup>1</sup>

<sup>1</sup> Max Planck Institutes Tübingen, Tübingen, Germany,  
{rakitsch,borgwardt,stegle}@tuebingen.mpg.de

<sup>2</sup> Microsoft Research, Los Angeles, California, USA  
lippert@microsoft.com

<sup>3</sup> Helsinki Institute for Information Technology HIIT, Department of Information and Computer Science,  
Aalto University, Helsinki, Finland,  
hande.topa@aalto.fi

<sup>4</sup> Zentrum fuer Bioinformatik, Eberhard Karls Universität Tübingen, Tübingen, Germany

<sup>5</sup> Helsinki Institute for Information Technology HIIT, Department of Computer Science,  
University of Helsinki, Helsinki, Finland,  
antti.honkela@hiit.fi

**Abstract.** RNA-Seq technology allows for studying the transcriptional state of the cell at an unprecedented level of detail. Beyond quantification of whole-gene expression, it is now possible to disentangle the abundance of individual alternatively spliced transcript isoforms of a gene. A central question is to understand the regulatory processes that lead to differences in relative abundance variation due to external and genetic factors. Here, we present a mixed model approach that allows for (i) *joint* analysis and genetic mapping of multiple transcript isoforms and (ii) mapping of *isoform-specific* effects. Central to our approach is to comprehensively model the causes of variation and correlation between transcript isoforms, including the genomic background and technical quantification uncertainty. As a result, our method allows to accurately test for shared as well as transcript-specific genetic regulation of transcript isoforms and achieves substantially improved calibration of these statistical tests. Experiments on genotype and RNA-Seq data from 126 human HapMap individuals demonstrate that our model can help to obtain a more fine-grained picture of the genetic basis of gene expression variation.

**Keywords:** RNA-Seq, eQTL, GWAS, alternative splicing

## 1 Introduction

**Motivation** Large-scale genotyping and expression profiling initiatives have fostered expression quantitative trait loci (eQTL) analyses, investigating the genetic component of the transcriptional state of the cell. In human, eQTL studies have revealed an abundance of genetic associations between proximal polymorphic loci and individual genes [27,28,20,18]. In model organisms, such as segregating yeast strains [4,24], mouse [22] or *Arabidopsis thaliana* [30], studies have uncovered a map of the genetic component of gene regulation, which is characterized by an interplay of local genetic associations in *cis* and distal genetic effects with *trans* mechanisms. The majority of existing studies of transcription were based on microarray technologies, which are limited to a coarse quantification of the total transcript abundance of *a-priori* known genes present in a sample.

Only recently, thanks to second generation sequencing techniques, deep transcriptome sequencing (RNA-Seq) has become viable even for population-scale analyses. Seminal work [18,20] has demonstrated the merits of using the digital readout provided by RNA-Seq technologies for genetic analyses, allowing for a more comprehensive dissection of the genetic component of transcriptional variability. RNA sequencing not only allows for an unbiased genome-wide quantification of transcription at an unprecedented resolution, it also enables the detection and quantification of alternative splice forms of genes [19,29]. Differences in the abundances of individual isoforms of a gene have been shown to play an important role for phenotypes in specific tissues. For example, genetic modifications that affect alternative splicing and post-transcriptional quality control mechanisms are postulated to be involved in a large proportion of genetic disorders, including the development of cancer [17,23,9,6,5].

By treating transcript isoform abundances as quantitative traits it is possible to perform eQTL studies on the level of individual isoforms. Such an approach allows for detecting polymorphic loci affecting the expression level of specific isoforms, for example via regulation of splice factors. While there have been previous attempts to tackle this problem at the level of individual exons [18,20], there is a lack of statistical models that allow for accurately dissecting the genetics of transcriptional gene-regulation at the level of individual transcript isoforms. Such methods face various challenges that need to be addressed in order to get meaningful results: First, the expression levels of individual transcript isoforms are highly correlated due to common transcriptional regulation as well as shared exons and splicing mechanisms between different isoforms. Second, limited identifiability of individual transcript isoforms from RNA-seq data causes correlation between isoform abundance estimates, which can confound the analysis if not accounted for. Finally, as in almost any genomic analyses of quantitative traits, there is the need to account for structure within the sample, for example due to shared ancestry or population clustering.

*Our goal in this article is to present a probabilistic model that allows for detecting (i) genetic loci that affect the overall expression level of multiple transcript isoforms and (ii) loci that act in an isoform-specific manner, while addressing the aforementioned challenges.*

**Related approaches** An alternative paradigm to what we present here are approaches that perform the analysis on the level of single exons instead of complete isoforms. Such analyses approaches have been used in the context of genetic perturbations [18,20] and other factors [2]. While such an approach imposes fewer assumptions on the processes that are involved, individual exon-specific regulation needs to be retrospectively combined to a transcript-level view. Furthermore, sensitive transcript expression quantification can borrow statistical strength across exons of the same transcript, potentially picking up more subtle variation (See also discussions in [8,12]).

Conceptually, our approach is related to multi-trait mixed models, which have undergone a long history of development (see e.g. [10,26,21]). Recently, these approaches have been successfully applied to genome-wide association studies of correlated traits and have been shown to increase power to detect pleiotropic SNP effects on two correlated traits as well as being able to detect environment-specific effects of SNPs on a phenotype measured under two different environments [13].

**Contributions of this paper** In this work we propose a mixed model to jointly map individual isoforms from alternatively spliced genes and to uncover genomic variants that affect *all* or *specific* transcript isoforms. By employing a state-of-the-art approach to infer expression levels of individual isoforms [8], we build a problem-specific background covariance that explains confounding genetic and technical correlation between transcript isoform abundance estimates of the same gene. We apply the

resulting mixed model to search for proximal *cis* associations in 126 HapMap individual where we find extensive regulation at gene level but also genetic effects that alter the expression abundance in a transcript-specific manner.

Besides a fine-grained picture of the genetic basis of gene expression, another key advantage of the proposed approach is that it is convenient to accurately measure statistical significance in our model.  $p$ -values for association can be computed analytically through likelihood ratio tests, without the need for expensive permutation experiments to assess statistical significance. By controlling for sources of confounding our model does not suffer from  $p$ -value inflation, as it is observed in simpler methods, and thus provides better control for the type-1 error rate.

## 2 Transcript expression inference from RNA-Seq

Transcript isoform expression can be only indirectly reconstructed from RNA-Seq datasets. Here, we build on BitSeq [8], a fully probabilistic approach to estimate transcript isoform expression levels from RNA-Seq data and a known transcript annotation.

Briefly, the generative model of the RNA-seq reads in each sample assumes a categorical assignment of reads to one of all genome-wide transcripts or a noise category according to  $p(I_j^n | \theta^n) = \text{Cat}(I_j^n | \{\theta_1^n, \dots, \theta_M^n\})$ , where  $\theta_t^n$  is the probability of observing a fragment of transcript isoform  $t$  in sample  $n$ , and  $M$  denotes the total number of annotated transcript isoforms. The likelihood of a read  $r_j$  given a transcript assignment,  $p(r_j | I_j = m) = p(r_s | \text{seq}_{mps})p(p | m)p(s | m)$ , is given by the probability of drawing a read from the specific position in the transcript isoform  $m$  while accounting for mismatches, position and sequence bias correction. Given an observed set of reads, we can apply Bayesian Markov chain Monte Carlo inference to obtain a full posterior probability distribution over the transcript isoform relative expression levels  $\theta_t^n$ . These are transformed to log-expression levels  $y_t^n = \log \theta_t^n$  for further analysis.

For each gene, we use the Monte Carlo estimates to obtain summary statistics that capture the mean expression level and covariation between the abundance estimates  $y_t^n$  of the  $T_g$  different transcripts of each gene  $g$  within each sample  $n$ . Modeling covariation is useful because transcript isoforms often share sequence which causes ambiguity in assigning reads to individual isoforms, and the transcript expression posteriors often show strong negative correlation (See also [8]). These transcript isoform means and between-isoform expression covariance summaries are then used for further modeling. This first and second moment of the variation for each gene is explained by a Gaussian over  $\mathbf{y}^n = \{y_1^n, \dots, y_{T_g}^n\}$ ,  $p(\mathbf{y} | \mathcal{D}) = \mathcal{N}(\mathbf{y} | \hat{\mathbf{y}}, \mathbf{C}^n)$ .

Along the same lines, the same Monte Carlo estimates can be use to estimate whole-gene expression, ignoring the transcript structure. For a specific gene  $g$  in sample  $n$ , expression levels can be estimated as  $y_g^n = \log \sum_{t \in T} \theta_t^n$ , where  $T$  is the set of all transcripts belonging to the gene  $g$ . These estimate yield RKPM estimates that are near-identical to standard counting-based methods.

## 3 A multi-isoform mixed model

To differentiate joint effects from isoform-specific regulation it is important to account for different sources of correlation between multiple transcript isoforms of the same gene. Here, we propose a multi-isoform variance component model that comprehensively accounts for different origins of correlation. First, multiple transcript isoforms of the same gene are correlated across samples because of genetic factors that affect multiple transcripts in the same way. In the absence of transcript-specific splicing regulation, this variation is dominated by a common gene expression component. Second, there is technical covariation between the isoforms quantified in the same sample (see Section 2).

In Section 3.1, we introduce a linear mixed model that accounts for both of these types of genetic and technical factors that induce correlation between transcripts. In Section 3.2, we show how the model can be used to test for common and specific genetic effects of a SNP (single nucleotide polymorphism).

### 3.1 Multi-isoform linear mixed model

For a specific gene (we drop the dependence on  $g$  to unclutter notation), let the  $N \cdot T$ -dimensional vector  $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_T]$  denote the vector of log expression levels for  $T$  different isoforms of a single gene measured in each of the  $N$  samples. For each  $t \in [1, \dots, T]$ , we model the  $N$ -dimensional vector  $\mathbf{y}_t = (y_t^1, \dots, y_t^N)$  holding all samples of isoform  $t$  by a linear model as follows:

$$\mathbf{y}_t = \underbrace{\mathbf{1} \cdot \mu_t}_{\text{Transcript mean}} + \underbrace{\mathbf{x}^* \cdot \beta^*}_{\text{SNP effect}} + \underbrace{\sum_{l \in \text{pop}} \mathbf{x}_l \cdot \beta_{l,t}}_{\text{population structure}} + \underbrace{\psi_t}_{\text{noise}}. \quad (1)$$

Here,  $\mu_t$  is the bias for transcript  $t$ ,  $\mathbf{x}^*$  is a SNP we would like to test for association, for example in the vicinity of the gene. The SNPs  $\{\mathbf{x}_l\}_{l \in \text{pop}}$  are genome-wide markers that capture population structure, excluding SNPs within a window of  $\pm 1$  mega base around the gene to avoid proximal contamination [16] when testing SNPs that lie in *cis*. The genetic *cis* effect  $\beta^*$  is modeled as a fixed effect. The population structure effects  $\{\beta_{l,t}\}_{l \in \text{pop}}$  are modeled as random. To allow for correlation between the effects of population structure across all isoforms, the  $T$ -dimensional vector,  $\beta_l = [\beta_{l,1}, \dots, \beta_{l,T}]$  is coupled across transcript isoforms by an identical multivariate normal distribution with covariance  $\mathbf{V}^{\text{pop}}$ , for each SNP  $\mathbf{x}_l$  used to represent population structure.

$$\beta_l \sim \mathcal{N}(\mathbf{0}, \mathbf{V}^{\text{pop}}); \quad \mathbf{V}^{\text{pop}} = \begin{bmatrix} v_{1,1}^{\text{pop}} & \dots & v_{1,T}^{\text{pop}} \\ \vdots & \ddots & \vdots \\ v_{1,T}^{\text{pop}} & \dots & v_{T,T}^{\text{pop}} \end{bmatrix} \quad (2)$$

The  $T(T+1)/2$  independent entries in  $\mathbf{V}^{\text{pop}}$  are inferred from the data within the model. For genes with a large number of transcript isoforms, the number of parameters may become prohibitive, in which shrinkage approaches such as Lasso regulation on the inverse covariance [7] could be used; see also Discussion.

The observational noise is independent across individuals, but correlates across isoforms within each individual due to technical variation. For every individual  $n$ , we let the  $T$ -dimensional vector  $\psi^n = [\psi_1^n, \dots, \psi_T^n]$  be the noise across all isoforms, which is distributed as

$$\psi^n \sim \mathcal{N}(\mathbf{0}, \text{diag}(\boldsymbol{\delta}) + \alpha^2 \mathbf{C}^n), \quad (3)$$

where  $\boldsymbol{\delta} = [\delta_1^2, \dots, \delta_T^2]$  is residual noise level of each isoform,  $\mathbf{C}^n$  denotes the technical correlation matrix between the isoforms due to quantification in sample  $n$  and  $\alpha^2$  is a common scaling parameter of the technical covariation. The technical correlation matrices  $\mathbf{C}^n$  are estimated *a-priori* by the BitSeq algorithm (see Section 2). The noise-variances  $[\delta_1^2, \dots, \delta_T^2]$ , and  $\alpha^2$  are estimated within the model.

*Marginal model* Integrating over the effects of population structure  $\beta_{l,t}$  for all  $l$  and  $t$  we get a joint marginal likelihood for all isoform levels of the gene:

$$\begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_T \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} \mathbf{1} \cdot \mu_1 \\ \vdots \\ \mathbf{1} \cdot \mu_T \end{bmatrix}; \mathbf{K}^{\text{pop}} + \mathbf{K}^{\text{noise}} \right), \quad (4)$$

where the covariance term consists of the genotype signal covariances

$$\mathbf{K}^{\text{pop}} = \begin{bmatrix} v_{1,1} \left( \sum_{l \in \text{pop}} \mathbf{x}_l \mathbf{x}_l^\top \right) & \dots & v_{1,T} \left( \sum_{l \in \text{pop}} \mathbf{x}_l \mathbf{x}_l^\top \right) \\ \vdots & \ddots & \vdots \\ v_{1,T} \left( \sum_{l \in \text{pop}} \mathbf{x}_l \mathbf{x}_l^\top \right) & \dots & v_{T,T} \left( \sum_{l \in \text{pop}} \mathbf{x}_l \mathbf{x}_l^\top \right) \end{bmatrix},$$

and the noise covariance

$$\mathbf{K}^{\text{noise}} = \alpha^2 \underbrace{\begin{bmatrix} \mathbf{D}_{1,1} & \dots & \mathbf{D}_{1,T} \\ \vdots & \ddots & \vdots \\ \mathbf{D}_{T,1} & \dots & \mathbf{D}_{T,T} \end{bmatrix}}_{\text{technical noise}} + \underbrace{\begin{bmatrix} \delta_1^2 \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \delta_T^2 \mathbf{I} \end{bmatrix}}_{\text{noise}},$$

where  $\mathbf{D}_{i,j} = \text{diag}([C_{i,j}^1, \dots, C_{i,j}^N])$ . Note that the noise covariance  $\mathbf{K}^{\text{noise}}$  is a sum of a diagonal matrix for observation noise and a chessboard pattern matrix consisting of diagonal matrices that couple different transcript within individual samples  $n$  due to technical noise. For an illustration of the effective covariance for a simple two-transcript gene, see Figure 1. Note that in the case of a gene with a single transcript, this variance component collapses to a standard linear mixed model with realized relationship matrix (see for example [11,15]).




Fig. 1: Illustration of the marginal covariance model for a gene with two transcript isoforms. The total covariance of the stacked transcript expression profiles are decomposed into a genetic component, technical covariance and transcript-specific observation noise.

We estimate the variance parameters of the model by maximizing the marginal likelihood (4) using constrained quasi-Newton optimization. To avoid prohibitive computations due to re-estimation of covariance parameters for different *cis* SNPs  $\mathbf{x}^*$ , we estimate all parameters only once for each gene, ignoring the fixed effect of the *cis* SNP. Analogous approximations, fitting the background covariance parameters on the null model, have previously been successfully applied for GWAS [11].

### 3.2 Mixed model to test for common and transcript isoform-specific genetic effects

We would like to test for the effect of individual *cis* SNPs that either have a common effect on all transcript isoforms or act in an isoform-specific manner. For this purpose, we use the parameters obtained from the multi-isoform variance components model Eqn. (4) in uni-variate testing. We construct a linear mixed model using different designs modeling the effect of  $\mathbf{x}^*$  as either common genetic effects across transcript isoforms or isoform-specific genetic effects:

$$\begin{bmatrix} \mathbf{y}_1 - \mathbf{1} \cdot \mu_1 \\ \vdots \\ \mathbf{y}_t - \mathbf{1} \cdot \mu_t \\ \vdots \\ \mathbf{y}_T - \mathbf{1} \cdot \mu_T \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} \mathbf{1} \\ \vdots \\ \mathbf{1} \\ \vdots \\ \mathbf{1} \end{bmatrix} \cdot b + \underbrace{\begin{bmatrix} \mathbf{x}^* \\ \vdots \\ \mathbf{x}^* \end{bmatrix}}_{\text{joint effect}} \cdot \beta_0 + \underbrace{\begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{x}^* \\ \vdots \\ \mathbf{0} \end{bmatrix}}_{\text{specific effect}} \cdot \beta^* ; \sigma_c^2 \underbrace{(\mathbf{K}^{\text{pop}} + \mathbf{K}^{\text{noise}})}_{\text{relatedness matrix}} + \sigma_e^2 \mathbf{I}_{N \cdot T} \right). \quad (5)$$

Here, the joint effect component is the identical SNP  $\mathbf{x}^*$  replicated over each transcript isoform. From this linear mixed model, different likelihood ratio tests can be carried out for each SNP in *cis*, using the computational tricks proposed in the FaST-LMM algorithm [15]. In the implementation, we specify the relatedness matrix to be the sum of  $\mathbf{K}^{\text{pop}}$  and  $\mathbf{K}^{\text{noise}}$  estimated on the null model. This approach enables us to obtain  $p$ -values from a  $\chi^2$  distribution. Related testing strategies, however focusing on pairs of quantitative traits in GWAS, have been previously used in [13].

*Gene-level association test* In order to test for *joint effects* of SNPs that act on transcription on a gene-level, namely  $\mathbf{x}^*$  that jointly regulate all transcript isoforms in the same direction, we fit the mixed model in Equation (5) only with the joint effect, discarding transcript-specific effects. This alternative model is compared to a null model by additionally dropping the joint effect and performing a one degree of freedom likelihood ratio test.

*Isoform-specific association test* In order to test for *specific effects* of SNPs, i.e. whether a SNP  $\mathbf{x}^*$  has an effect that acts specifically on a transcript isoform  $t$ , we fit the full mixed model in Equation (5), placing the specific effect only on the abundance of this isoform  $\mathbf{y}_t$ . This alternative model is compared to a null model that drops the specific effect. By conditioning on the joint effects we ensure that we only retrieve effects that act differently between the transcript isoforms. This results in a likelihood ratio test with one degree of freedom.

*Combined association test* Finally, we consider a multi-transcript test that assesses whether there is *any* association with a transcript isoform  $t$ , either by jointly regulating all isoforms or by acting specifically on the isoform  $t$ . The alternative model is the same as when testing for isoform-specific tests. However, in this case we drop both, the joint effect as well as the specific effect from the null model. In this case the likelihood ratio test has two degrees of freedom.

In the following we will denote these three testing strategies when using the multi transcript covariance structure fitted on the null model (Section 3.1) as multi-isoform mixed model (MIMM). We also consider alternative methods for the purpose of comparison, for example without the multi trait structure and the technical noise covariance.

## 4 *cis* QTL mapping of transcript isoform levels in HapMap populations

We applied the multi-isoform mixed model (MIMM) and alternative methods to publicly available RNA-Seq data of human HapMap samples [1], combining the data from Pickrell et. al [20] and Montgomery et. al [18]. Altogether, there were 59 samples from the CEU population and 74 samples from the YRI HapMap population that could be mapped to unique HapMap identifiers. SNPs were filtered for minimum allele frequency 0.05 and we discarded variants with missing genotyping information in more than 5% of the samples. Finally, we also discarded all samples with more than 5% missing SNPs. After filtering, our dataset comprised of a total of 126 individuals and 996,158 SNPs.

Transcript quantification was done using BitSeq (Section 2), where we used the GRCh37.p68 reference transcriptome from Ensembl and mapped the RNA-Seq reads with bowtie [14]. As the focus of this work resides on transcript isoform-specific effects, we only considered genes with at least two transcript isoforms. We used the marginal uncertainty estimates from BitSeq (Section 2) to filter out transcript isoforms that were consistently not expressed or difficult to quantify in 90% of the samples or more ( $z$ -score cutoff 1.5). Because of the small sample size in these datasets, genes with more than 4 active isoforms were discarded, ensuring that the comparison of methods was not compromised by sample size. This final dataset consisted of 5,954 genes, comprising of a total number of 16,340 individual transcripts.

Statistical tests for proximal *cis* effects were carried out in a 1 Mb window upstream or downstream of each gene start and stop. In both the MIMM and ordinary mixed models, we considered for comparison, we cut out this region for building the multi-isoform relationship matrix  $\mathbf{K}^{\text{pop}}$  to avoid possible loss of power due to proximal contamination [16]. We also included an indicator variable distinguishing between the YRI and CEU population to account for this low-rank component of population structure. Estimates of a local false discovery rate for each gene were obtained from the Benjamini & Hochberg [3] procedure to estimate  $q$ -values.

#### 4.1 Testing for general regulatory effects on whole gene expression and transcript isoform levels

First, we wanted to relate results obtained by transcript-based modeling, as proposed here, with associations from an eQTL scan on whole gene expression, as considered in previous analyses [18,20]. To this end, we applied the multi-isoform mixed model (MIMM) and performed the combined association test to find genetic effects that are either specific to individual transcript isoforms or act on the gene-level. We compared these results with a linear mixed model (LMM), carrying out standard eQTL tests on whole gene expression levels that were also estimates from BitSeq in a manner consistent with the transcript level estimates (Section 2). Many but not all genes that had at least one significant *cis* association by one of the methods ( $q$ -value threshold of 5%) were also detected by the other approach (see Figure 2 (a)). This result suggests that the combined association test on a transcript level is capable of retrieving associations that correspond to established eQTLs but also other signals. As expected, both models yielded associations with a strong enrichment for lying in close proximity to the gene start (Figures 2 (b) and (c)), suggesting that many of these hits are genuine QTLs. We also carried out genome-wide scans for a random selection of 49 of genes to investigate calibration of  $p$ -values, finding that both methods yielded calibrated  $p$ -values (genomic control:  $\lambda = 1.02$  MIMM,  $\lambda = 1.03$  LMM). Summary results are shown in Table 1).




Fig. 2: Comparison of significant results ( $q$ -value threshold 5%) of the MIMM model, testing for combined associations, and a standard linear mixed model applied to whole gene expression levels. (a) Mutual overlap of associations. (b-c) Density of the distance to the gene for *cis* associations.

#### 4.2 Dissecting genetic associations specific to individual transcript isoforms

Next, we strived to dissect the set of associations retrieved by the combined association test into associations that are either common to all transcripts, i.e. occur at a gene-level, and transcript-specific genetic effects. To this end, we used the MIMM and additionally carried out gene-level association tests and isoform-specific association tests in the same local *cis* regions as before. The overlap of results obtained from the three alternative testing strategies are shown in Figure 3. The combined association test retrieved the greatest number of genes with at least one association, overlapping with many of the hits from either the gene-level test or the transcript level test. Notably, genes with an associations found by the gene-level test and the transcript-specific test were almost completely disjoint, suggesting that

	# significant genes			# significant isoforms			# significant SNPs			Genomic Control ( $\lambda$ )		
	LIN	LMM	MIMM	LIN	LMM	MIMM	LIN	LMM	MIMM	LIN	LMM	MIMM
<i>Gene expression tests</i>												
general genetic effect	612	591	-	-	-	-	6417	6011	-	1.03	1.03	-
<i>Transcript expression tests</i>												
combined genetic effect	4,085	3,608	498	11,302	9,000	1,016	992,017	975,583	10,521	1.37	1.51	1.02
isoform-specific effects	2,497	3,498	429	6,110	8,430	707	991,041	1,226,504	4,315	1.12	1.73	1.00
gene-level effects	2,581	453	441				88,494	18,546	4,516	1.42	1.02	1.02

Table 1: Summary statistics of the associations retrieved by alternative methods on different expression traits. Top row: Gene expression test, carried out on whole-gene expression estimates ignoring transcript structure. Bottom row: Transcript expression tests, carried out on all transcripts within a gene jointly using alternative testing strategies.




Fig. 3: Overlap of associations retrieved by the multi-isoform mixed model when considering alternative testing strategies.

within individual genes there is a strong tendency for either gene-level regulation or transcript-level regulation.

Perhaps the most relevant use of multi transcript-isoform models is the association analysis of isoform-specific regulatory effects. Figure 4 (f) shows the distance of the 707 isoform-specific associations relative to the transcript start site, showing a similar enrichment as observed for gene-level associations (Section 4.1). Again, we considered a selection of transcripts and carried out genome-wide scans to investigate calibration of p-values (Figure 4 (c)). This apparent calibration of tests ( $\lambda = 1.00$ ) when using the MIMM is crucially linked to the fitted multi trait covariance structure, accounting for transcript correlation and different sources of noise (Section 3.1). To explore the impact of this modeling, we also considered simpler alternative covariance structure, either as used in a standard linear mixed model (LMM) or a linear model without any background covariance (LIN). At a  $q$ -value cutoff of 5%, the linear model and the linear mixed model retrieved a greater number of genes with at least one transcript-specific association (Table 1). However, these absolute numbers need to be put in context with the statistical calibration of different methods: Figure 4a-c show genome-wide QQ-plots for a random selection of 121 transcripts (49 genes) tested for association by each of the considered methods. The excess of small  $p$ -values retrieved by the linear model ( $\lambda = 1.12$ ) and the linear mixed model ( $\lambda = 1.73$ ) suggests that the large number of findings is due to severely inflated the test statistics and hence the majority of these hits are likely to be false positives rather than actual true associations. This hypothesis is further supported by the lack of enrichment of associations in proximity to the transcription start site (Figure 4d-f, where only the multi-isoform mixed model yielded hits consistent with the hypothesis that true regulatory variants are likely to be within the gene body. Similar deficiencies were also found when testing for joint genetic effects (See Figure 5).

We also investigated the distribution of the number of isoform-specific associations per gene. The majority of genes had either a single transcripts that was under *cis* genetic regulation (38.00%) or had



two regulatory events (59.67%). Only a small minority of genes had more complex regulation involving more than two transcripts (2.33%).

Two example genes, one with a joint genetic effect and a transcript-specific effect are shown in Figure 6. This illustrations demonstrates how different testing approaches provide a finer-grained picture of the genetic regulation of gene expression, both at the gene- and at the transcript level.

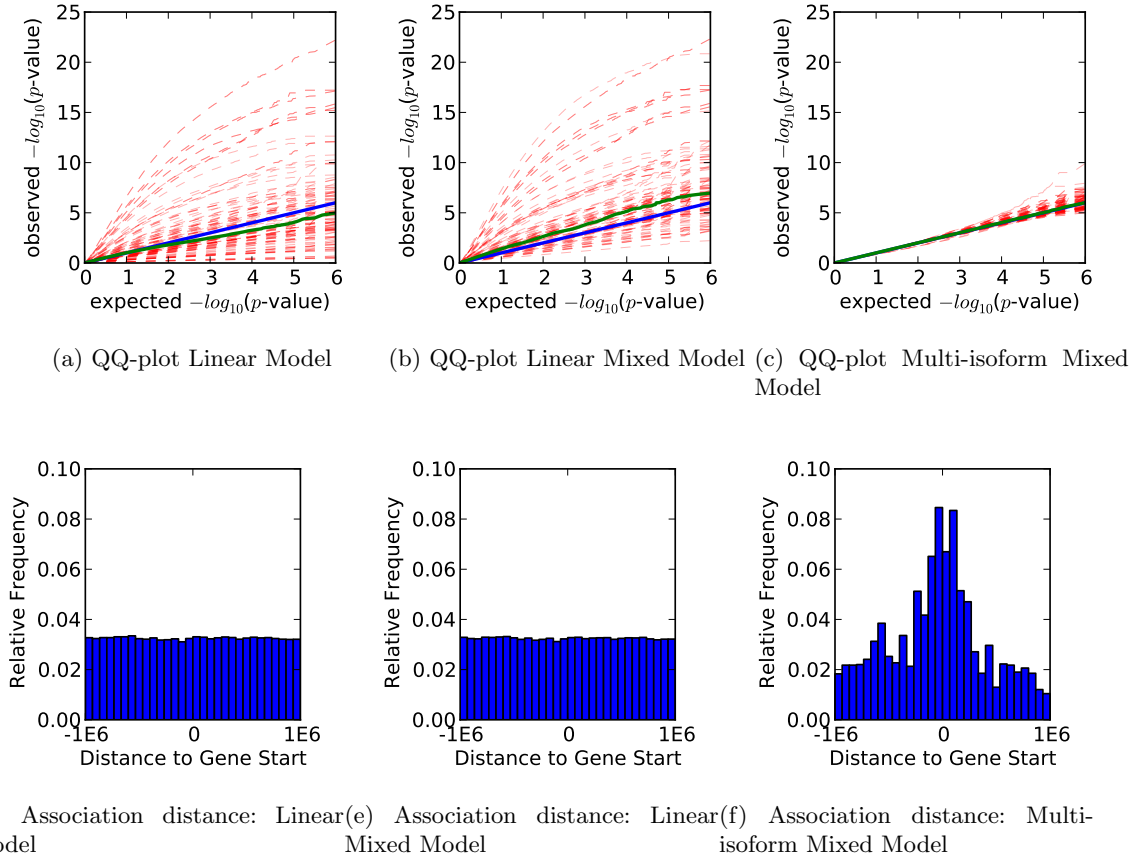


Fig. 4: Comparison of alternative methods to detect transcript-specific associations in proximal *cis* regions. (a)-(c): Quantile-quantile plots for genome-wide scans of a selection of 121 traits using alternative methods. (d)-(f): Histogram of the distance of the most associated variant to the transcription start cite for significant hits (FDR 5%). The multi-isoform mixed model achieved the best calibration of test statistics and retrieved *cis* associations with an enrichment near the transcription start site.

## 5 Conclusions

Transcriptional and post-transcriptional regulation are complex biological processes, both of which are under genetic control. Here, we have introduced a mixed model approach to test for genetic effects that either act on all transcripts *jointly* or alter expression levels of *specific* transcript isoforms while others remain unaltered.

In a proof of concept application to RNA-Seq profiles from 126 HapMap samples, we have demonstrated substantial benefits of this approach compared to established analysis strategies. First, we have

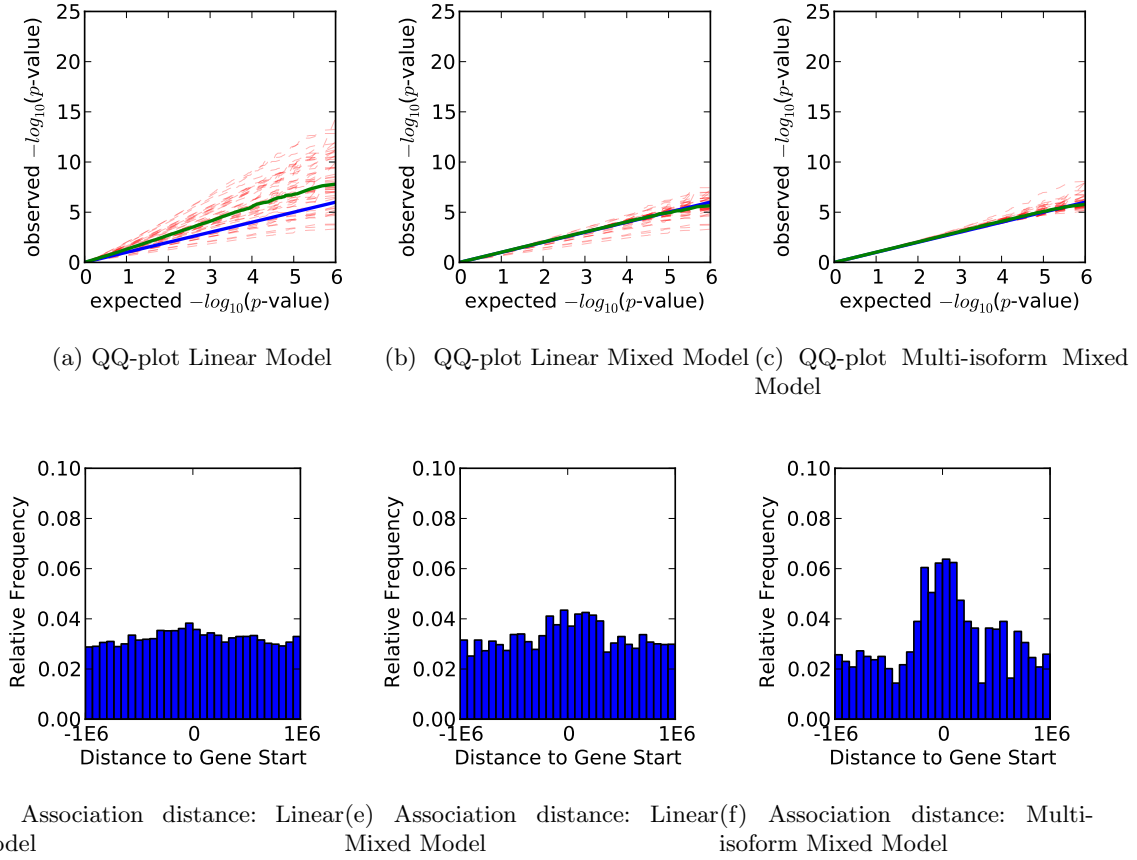


Fig. 5: Comparison of alternative methods to detect joint associations in proximal *cis* regions. **(a)-(c)**: Quantile-quantile plots for genome-wide scans of a selection of 121 traits using alternative methods. **(d)-(f)**: Histogram of the distance of the most associated variant to the transcription start cite for significant hits (FDR 5%). The linear mixed model and the multi-isoform mixed model showed improved calibration over the linear model, while the multi-isoform mixed model detected most *cis* associations near the transcription start site.

shown that accounting for covariation between transcripts within the same gene, both due to genetic and technical factors, greatly improves statistical calibration of  $p$ -values for different types of association tests. The mixed model covariance structure we propose is tailored to the analysis of estimates of multiple transcript isoform abundances. While general multi-trait mixed modeling has a long history for joint modeling of multiple traits (e.g. [10,26,21,13]), we are not aware of any practical application to larger sets than pairs of traits.

Second, we have shown how different testing strategies on the level of transcript isoforms yield more detailed mechanistic insights compared to existing approaches that operate on a gene level, while not loosing power. Using statistical tests that correspond to specific regulatory hypotheses, we were able to dissect this catalog of general *cis* associations into common genetic effects operating *across transcripts* and genetic effects that act in a *transcript-specific* manner. While common regulatory effects are most frequent, we have found evidence for transcript-specific regulation in 7.21% of genes; in some instances both types of effect were even found within the same gene.

In conclusion, we have shown how the combination of multi-trait mixed models and probabilistic transcript quantification is able to uncover novel biological insights while providing statistically robust estimates. Currently available RNA-Seq datasets are limited, both in sample sizes and in terms of read lengths for reliable isoform quantification. Because of these limitations, we focused on medium-

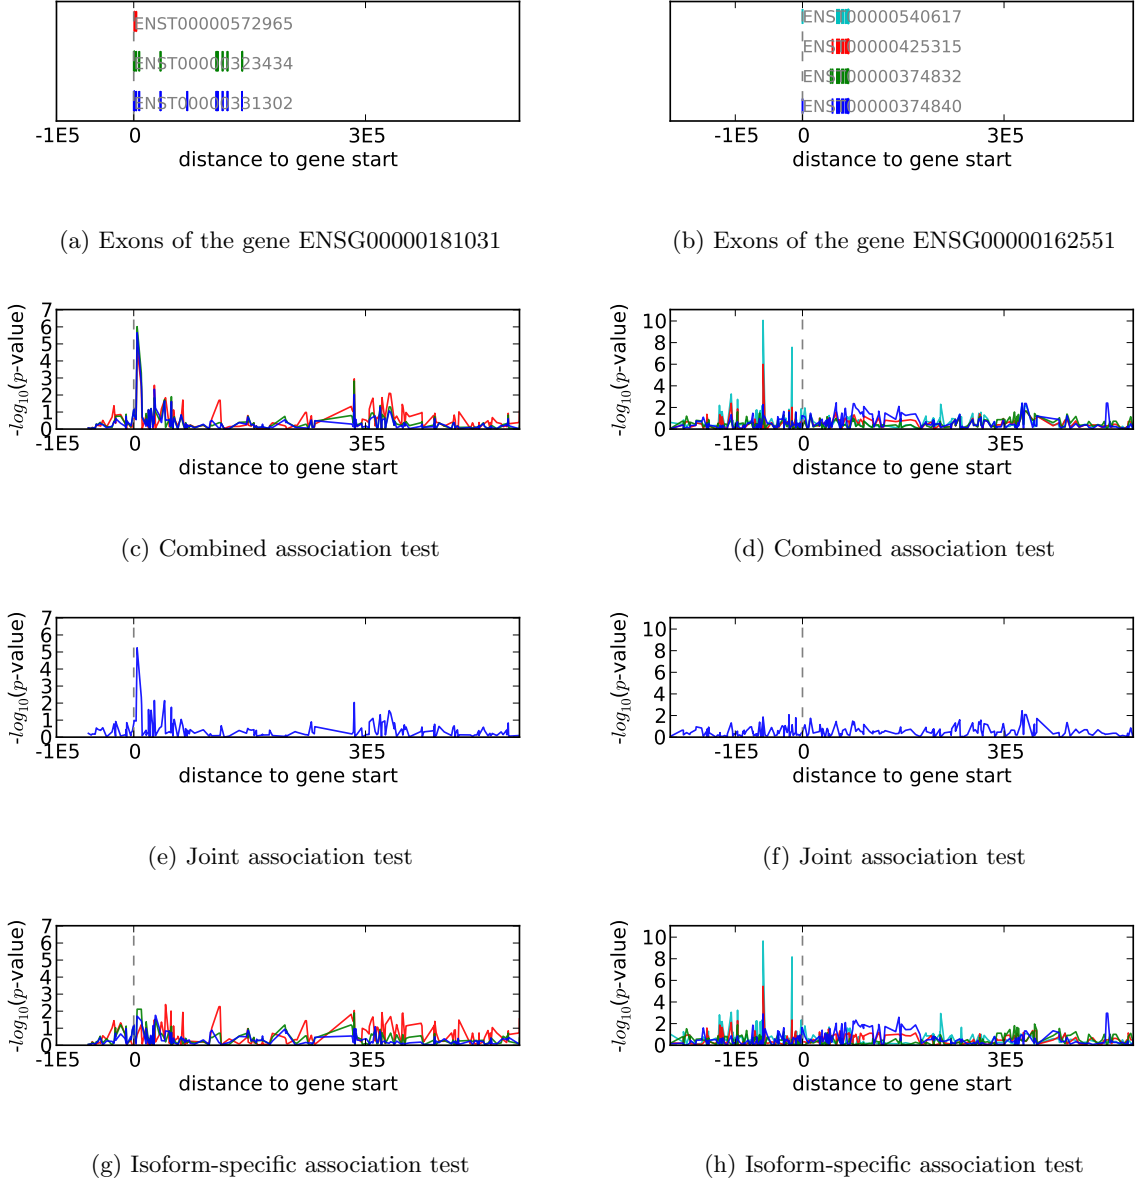


Fig. 6: Left side: Gene *rph3al* (ENSG00000181031) has an joint association for all three transcripts ENST00000331302(blue), ENST00000323434(green), ENST00000572965(red). Right side: Gene *alpl* (ENSG00000162551) has two distinct transcript-specific association for the transcript ENST00000540617 (cyan). The first one is shared with the transcript ENST00000425315 (red), while the second is not. The other two transcripts show no significant association.

complexity genes with 2–4 transcripts. More complex transcript structures may be fit by employing shrinkage priors, regularizing the effective number of parameters in the model [7,25]. In the future, larger datasets of better quality will render these tasks easier and hence models as the one proposed here will gain wide-spread applicability to many of the RNA-Seq eQTL studies to come.

## References

1. Altshuler, D., Gibbs, R., Peltonen, L., Dermitzakis, E., Schaffner, S., Yu, F., Bonnen, P., De Bakker, P., Deloukas, P., Gabriel, S., et al.: Integrating common and rare genetic variation in diverse human populations. *Nature* 467(7311), 52 (2010)
2. Anders, S., Reyes, A., Huber, W.: Detecting differential usage of exons from rna-seq data. *Genome Research* (2012)
3. Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 289–300 (1995)
4. Brem, R.B., Yvert, G., Clinton, R., Kruglyak, L.: Genetic dissection of transcriptional regulation in budding yeast. *Science (New York, N.Y.)* 296(5568), 752–5 (Apr 2002)
5. Danckwardt, S., Neu-Yilik, G., Thermann, R., Frede, U., Hentze, M., Kulozik, A.: Abnormally spliced  $\beta$ -globin mrnas: a single point mutation generates transcripts sensitive and insensitive to nonsense-mediated mrna decay. *Blood* 99(5), 1811–1816 (2002)
6. Fackenthal, J., Godley, L.: Aberrant rna splicing and its functional consequences in cancer cells. *Disease models & mechanisms* 1(1), 37–42 (2008)
7. Friedman, J., Hastie, T., Tibshirani, R.: Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 9(3), 432–441 (2008)
8. Glaus, P., Honkela, A., Rattray, M.: Identifying differentially expressed transcripts from RNA-seq data with biological variation. *Bioinformatics* 28(13), 1721–1728 (Jul 2012)
9. He, C., Zhou, F., Zuo, Z., Cheng, H., Zhou, R.: A global view of cancer-specific transcript variants by subtractive transcriptome-wide analysis. *PLoS One* 4(3), e4732 (2009)
10. Henderson, C., Quaas, R.: Multiple trait evaluation using relatives’ records. *Journal of Animal Science* 43(6), 1188–1197 (1976)
11. Kang, H.M., Sul, J.H., Service, S.K., Zaitlen, N.A., Kong, S.Y., Freimer, N.B., Sabatti, C., Eskin, E.: Variance component model to account for sample structure in genome-wide association studies. *Nature genetics* 42(4), 348–54 (April 2010)
12. Katz, Y., Wang, E., Airoidi, E., Burge, C.: Analysis and design of rna sequencing experiments for identifying isoform regulation. *Nature methods* 7(12), 1009–1015 (2010)
13. Korte, A., Vilhjálmsson, B., Segura, V., Platt, A., Long, Q., Nordborg, M.: A mixed-model approach for genome-wide association studies of correlated traits in structured populations. *Nature Genetics* 44(9), 1066–1071 (2012)
14. Langmead, B., Trapnell, C., Pop, M., Salzberg, S.L.: Ultrafast and memory-efficient alignment of short DNA sequences to the human genome. *Genome biology* 10(3), R25 (2009)
15. Lippert, C., Listgarten, J., Liu, Y., Kadie, C., Davidson, R., Heckerman, D.: FaST linear mixed models for genome-wide association studies. *Nature Methods* 8, 833–835 (2011)
16. Listgarten, J., Lippert, C., Kadie, C., Davidson, R., Eskin, E., Heckerman, D.: Improved linear mixed models for genome-wide association studies. *Nature Methods* 9(6), 525–526 (2012)
17. Matlin, A., Clark, F., Smith, C.: Understanding alternative splicing: towards a cellular code. *Nature Reviews Molecular Cell Biology* 6(5), 386–398 (2005)
18. Montgomery, S., Sammeth, M., Gutierrez-Arcelus, M., Lach, R., Ingle, C., Nisbett, J., Guigo, R., Dermitzakis, E.: Transcriptome genetics using second generation sequencing in a caucasian population. *Nature* 464(7289), 773–777 (2010)
19. Mortazavi, A., Williams, B.A., McCue, K., Schaeffer, L., Wold, B.: Mapping and quantifying mammalian transcriptomes by RNA-Seq. *Nature Methods* 5(7), 621–628 (Jul 2008)
20. Pickrell, J., Marioni, J., Pai, A., Degner, J., Engelhardt, B., Nkadori, E., Veyrieras, J., Stephens, M., Gilad, Y., Pritchard, J.: Understanding mechanisms underlying human gene expression variation with rna sequencing. *Nature* 464(7289), 768–772 (2010)
21. Price, A., Helgason, A., Thorleifsson, G., McCarroll, S., Kong, A., Stefansson, K.: Single-tissue and cross-tissue heritability of gene expression via identity-by-descent in related or unrelated individuals. *PLoS Genetics* 7(2), e1001317 (2011)
22. Schadt, E., Lamb, J., Yang, X., Zhu, J., Edwards, S., GuhaThakurta, D., Sieberts, S., Monks, S., Reitman, M., Zhang, C., et al.: An integrative genomics approach to infer causal associations between gene expression and disease. *Nature genetics* 37(7), 710–717 (2005)
23. Skotheim, R., Nees, M.: Alternative splicing in cancer: noise, functional, or systematic? *The international journal of biochemistry & cell biology* 39(7), 1432–1449 (2007)
24. Smith, E.N., Kruglyak, L.: Gene-environment interaction in yeast gene expression. *PLoS Biology* 6(4), e83 (Apr 2008)

25. Stegle, O., Lippert, C., Mooij, J., Lawrence, N., Borgwardt, K.: Efficient inference in matrix-variate gaussian models with iid observation noise. In: *Advances in Neural Information Processing Systems 24*. pp. 630–638 (1 2012)
26. Stich, B., Piepho, H., Schulz, B., Melchinger, A.: Multi-trait association mapping in sugar beet (*beta vulgaris* l.). *TAG Theoretical and Applied Genetics* 117(6), 947–954 (2008)
27. Stranger, B.E., Nica, A.C., Forrest, M.S., Dimas, A., Bird, C.P., Beazley, C., Ingle, C.E., Dunning, M., Flicek, P., Koller, D., Montgomery, S., Tavaré, S., Deloukas, P., Dermitzakis, E.T.: Population genomics of human gene expression. *Nature Genetics* 39(10), 1217–24 (2007)
28. Stranger, B., Montgomery, S., Dimas, A., Parts, L., Stegle, O., Ingle, C., Sekowska, M., Smith, G., Evans, D., Gutierrez-Arcelus, M., et al.: Patterns of cis regulatory variation in diverse human populations. *PLoS genetics* 8(4), e1002639 (2012)
29. Wang, Z., Gerstein, M., Snyder, M.: RNA-Seq: a revolutionary tool for transcriptomics. *Nature Reviews Genetics* 10(1), 57–63 (2009)
30. West, M., Kim, K., Kliebenstein, D., Van Leeuwen, H., Michelmore, R., Doerge, R., Clair, D.: Global eqtl mapping reveals the complex genetic architecture of transcript-level variation in arabidopsis. *Genetics* 175(3), 1441–1450 (2007)